

An Empirical Study of Variables Acceptance Sampling: Methods, Implementation, Testing, and Recommendations

K. Preston White, Jr.

Professor of Systems and Information Engineering, University of Virginia

Visiting Professor of Operations Research, Naval Postgraduate School

Member, NASA Engineering Statistics Team

Kenneth L. Johnson

Statistics and Trending

NASA/ NESC Systems Engineering Office

NASA Statistical Engineering Symposium

04 May 2011

Overview

- Acceptance sampling/Sampling plans
- Motivation
- Components of a probabilistic requirement
- Current NASA best practice (ASA)
- A potentially more efficient practice (ASV)
- Research plan, summary results, literature review
- Operating characteristic
- Derivation of variables sampling plans
- ASV sampling plan calculators
- Empirical testing and results
- Tests of the fundamental assumption (near normal, near exponential skew, and modest skew)
- Procedure for selecting a sampling plan (flow diagram)
- Summary/Contributions

Acceptance sampling

- One of the oldest problems in quality engineering is to assess the acceptability of items that a customer receives from a producer.
- Acceptance sampling is an alternative to 100% inspection applied when inspection is destructive, or when the time and/or cost of 100% inspection are unwarranted or prohibitive.
- Based on inspection of the sample, the customer decides whether to accept or reject the entire lot, or to continue sampling.
- There are standards (MIL, ANSI, and ISO) pertaining to acceptance sampling.

Sampling plans

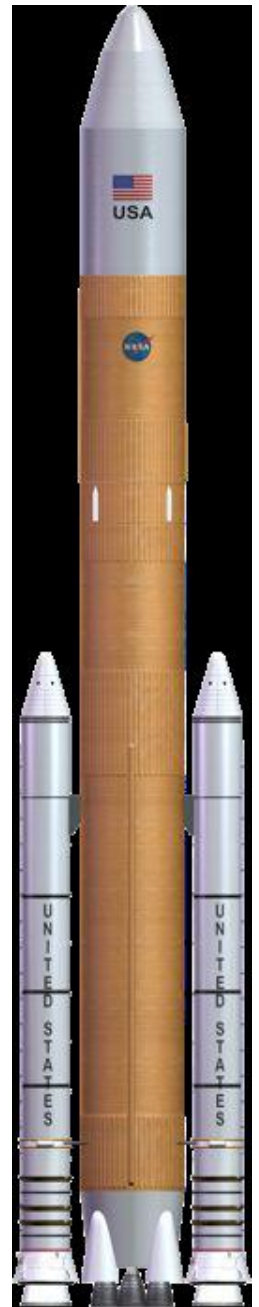
- A sampling plan is the pair (n, c) or (n, k) , where n is the minimum sample size, i.e., the minimum number of observations required to verify statistically the requirement.
- For discrete random variables, the constant c is the maximum number of nonconforming observations supporting the determination that a lot is acceptable.
- For continuous random variables, constant multiplier k is the minimum distance (in standard deviations) between the sample mean and the required limit supporting the determination that a lot is acceptable.

Motivation

Our interest in acceptance sampling arose in an analogous sampling experiment--the need to verify level-two design requirements for Cx “by analysis” using Monte Carlo simulation.

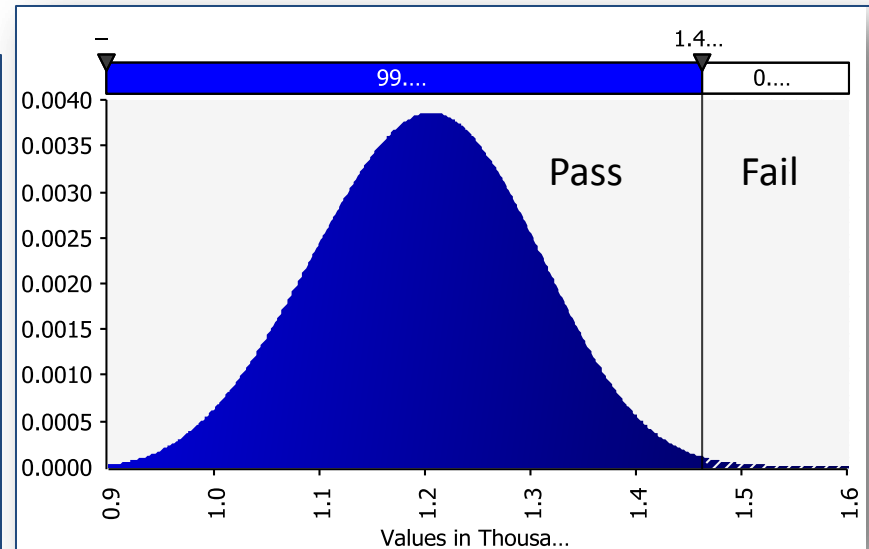
[CA0049-PO] The CaLV [Cargo Launch Vehicle] shall launch LSAM [Lunar Surface Access Module] from the launch site to the Earth Rendezvous Orbit (ERO) for Lunar Sortie Crew and Lunar Outpost Crew missions

The delivery of the LSAM from the launch site to the ERO shall be verified by analysis. The analysis shall be performed using NASA-accredited digital flight simulations. The analysis shall include Monte Carlo dispersions on mass properties, engine performance, GN&C parameters and environmental parameters. The verification shall be considered successful when the analysis results show that there is a $[\rho]$ probability with a $[100(1-\beta)\%]$ confidence that the LSAM reaches ERO.



Components of a probabilistic design requirement

- Condition (I)
conformance indicator (typically a limit on the value of an output variable)
- Reliability (ρ)
minimum probability of achieving the condition
- Consumer's risk (β)
maximum probability of accepting a nonconforming design
- Producer's risk (α)
maximum probability of rejecting a conforming design

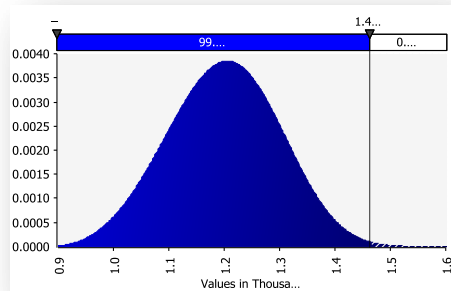


Consider the (true but unknown) parent distribution of an output variable X .

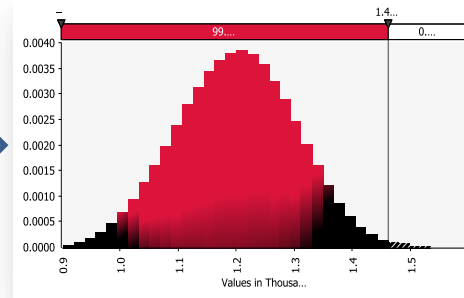
We can see that this output meets the condition $X \leq 1463$ with reliability $\rho = 0.997$.

If we knew the parent distribution a priori, there would be no sampling error and the risks would be $\beta = \alpha = 0$.

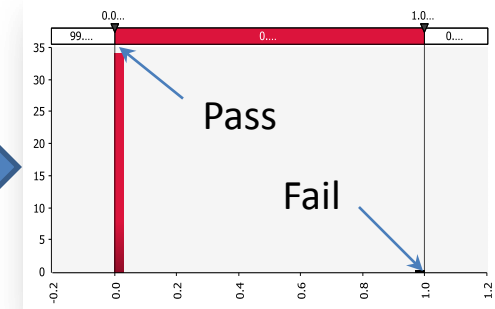
Current NASA best practice for requirements verification using Monte Carlo



Sample



Count



Sample the parent distribution using Monte Carlo simulation

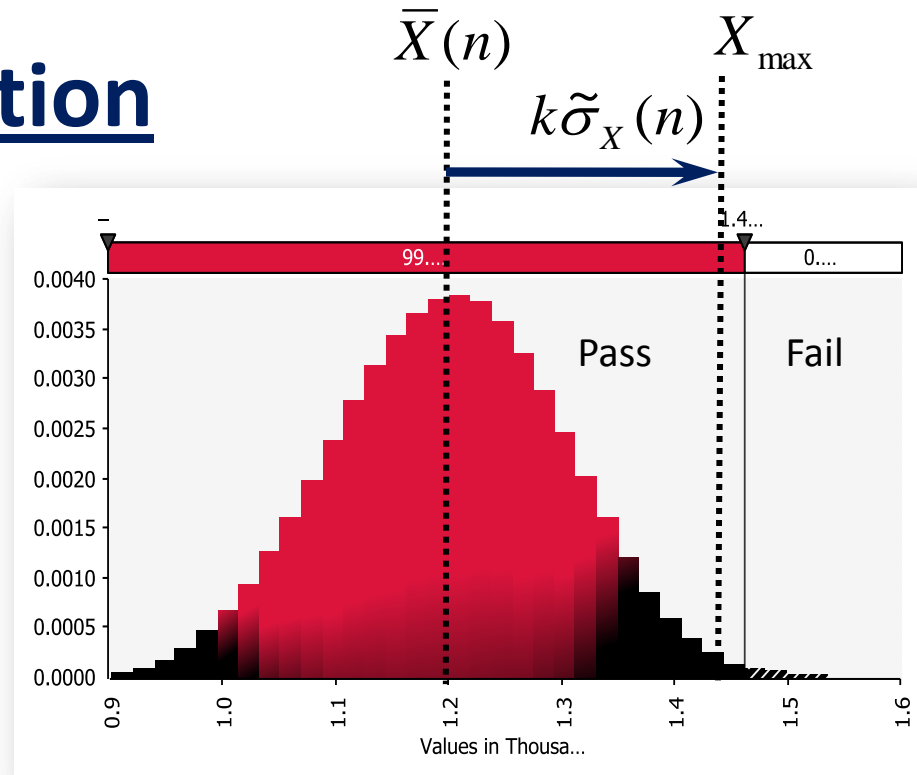
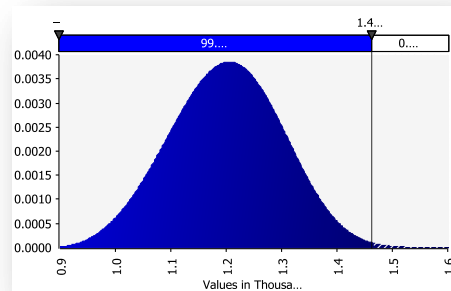
Count the number of simulation trials in which the output fails the condition

The current best practice employs **attributes acceptance sampling (ASA)**. For the required reliability and consumer's risk, the sampling plan specifies **number of trials (n)** and the **maximum number of failures permitted (c)** to substantiate the validity of the design.

Advantage: Plans are exact and can be determined a priori. (Nonparametric--by definition, the distribution of the count is $\text{Binomial}(n, p)$, where p is the true reliability.)

Disadvantage: Plans require large samples for high confidence in highly reliable designs (the pass/fail count ignores "by how much").

A more economical approach to verification



The current project seeks a best practice employing **variables acceptance sampling (ASV)**. For the required reliability and consumer's risk, the sampling plan specifies **number of trials** (n) and the **minimum multiplier** (k) to substantiate the validity of the design.

Advantage: ASV plans typically require fewer trials than ASA plans (but not always).

Disadvantages: Software for plan generation is unavailable; procedures/assumptions reported in the academic literature appear to be largely untested.

Research plan and summary results

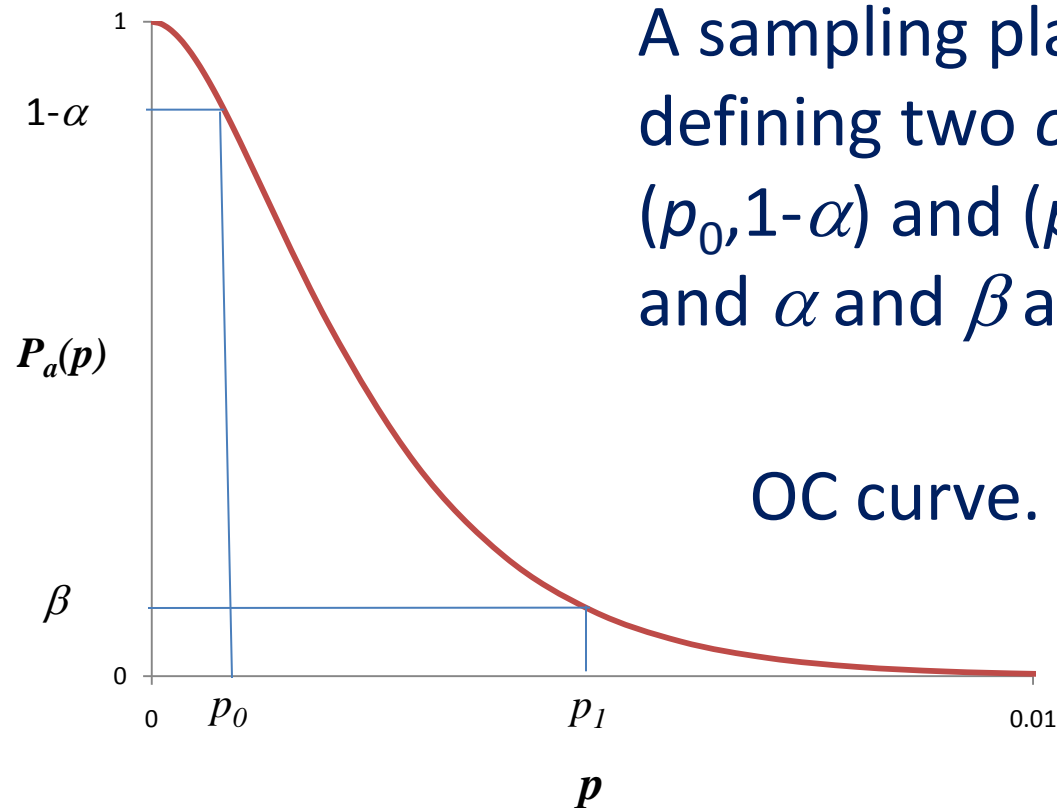
Software search	Off-the-shelf plan calculators (commercial or otherwise) were found <u>only for normal variates</u> .
Literature search	Plans for 5 additional variates were found in the academic literature (Exponential, Gamma, Weibull, Inverse Gaussian, Poisson, Burr).
Implementation	Calculators were implemented in Excel for Binomial, Normal, Exponential, Gamma, Weibull, Inverse Gaussian, and Poisson. (Burr not attempted.) Verified against published examples. Plans typically, <u>but not invariably</u> , smaller than corresponding ASA plans.
Empirical Testing	Monte Carlo simulation applied to test plans derived for typical (Constellation-like) OC from all seven calculators. All were validated, <u>except</u> for Inverse Gaussian. Error in the published IG derivation discovered.
Application issues	Fundamental assumption that the distributional form can be determined uniquely from sample data tested using Monte Carlo. Assumption <u>not</u> substantiated for typical OC. Conservative protocol developed for selecting plans to use in practice.

Literature

Variable	Source	Implemented	Validated
Binomial	Multiple sources	✓	✓
Normal	Multiple sources	✓	✓
Gamma	K. Takagi (1972) "On designing unknown-sigma sampling plans on a wide class of non-normal distributions," <i>Technometrics</i> 14(3)669-678.	✓	✓
Weibull	K. Takagi (1972) "On designing unknown-sigma sampling plans on a wide class of non-normal distributions," <i>Technometrics</i> 14(3)669-678.	✓	✓
Exponential	W. C. Guenther (1977), <i>Sampling Inspection in Statistical Quality Control</i> , Macmillan, New York.	✓	✓
Poisson	W. C. Guenther (1977), <i>Sampling Inspection in Statistical Quality Control</i> , Macmillan, New York.	✓	✓
Inverse Gaussian	M. S. Aminzadeh (1996), "Inverse-Gaussian Acceptance Sampling Plans by Variables," <i>Communications in Statistics--Theory and Methods</i> 25(5)923-935.	✓	×
Burr	K. Takagi (1972) "On designing unknown-sigma sampling plans on a wide class of non-normal distributions," <i>Technometrics</i> 14(3)669-678.	No	No

Operating Characteristic

Every sampling plan has an *operating characteristic* (OC) which defines the probability of accepting a population $P_a(p)$ for every value of the failure probability $p \in [0,1]$.



A sampling plan is derived by defining two *operating points*, $(p_0, 1-\alpha)$ and (p_1, β) , where $p_0 < p_1$ and α and β are small probabilities.

Derivation of variables plans

- The underlying problem can be framed as an hypothesis test for which we intend to enforce both significance and power requirements.

- The null and alternate hypotheses are

$$\mathbf{H}_0: p = p_0 \text{ and } \mathbf{H}_1: p = p_1 > p_0$$

- Under \mathbf{H}_0 we accept the population as conforming and under \mathbf{H}_1 we reject the population as nonconforming.

- The inequalities

$$P_a(p_0) \geq 1-\alpha \text{ and } P_a(p_1) \leq b$$

establish the significance and power of the test.

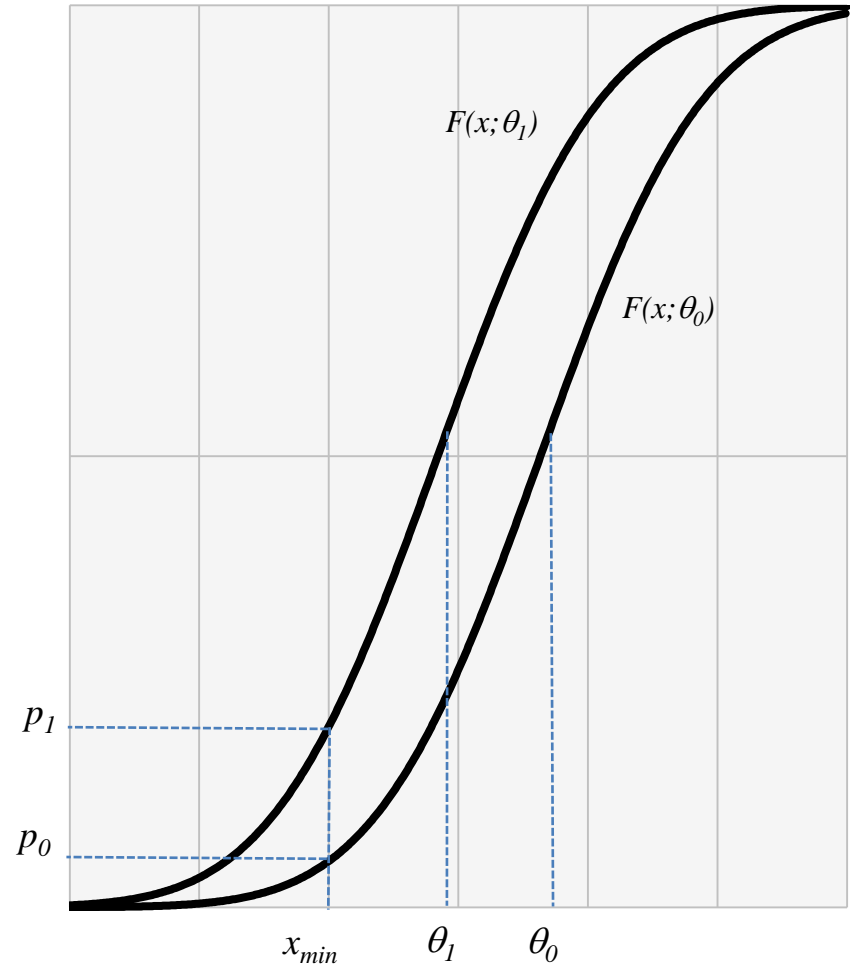
Derivation of variables plans

With the form of the distribution $F(x; \theta)$ known, the null and alternate hypotheses are equivalent to

$$H_0: \theta = \theta_0$$

$$H_1: \theta = \theta_1 > \theta_0$$

as shown for a required lower bound x_{\min} .

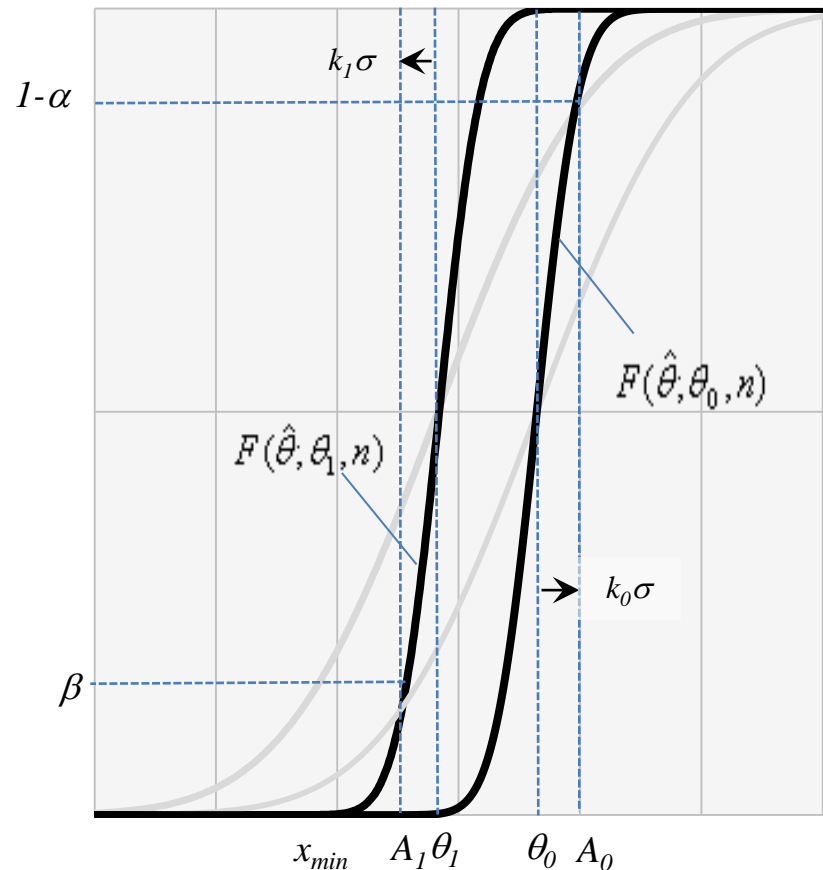


Derivation of variables plans

The power requirements are applied to the sampling distribution

$$F(\hat{\theta}; \theta, n)$$

to determine the acceptance limit A , required sample size n , and multiplier k .



Dashboard for the Weibull Calculator

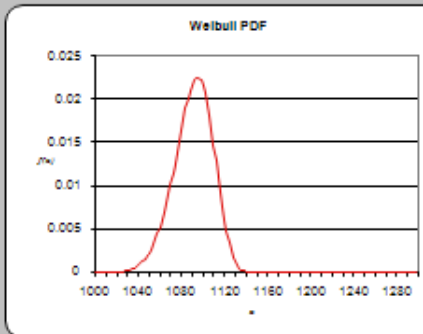
Weibull(v, λ, γ) Acceptance Sampling Plans

rev KPw 18 April 2009

ENTER the parameters of the Weibull(v, λ, γ) distribution in the fields below. The statistics for the Weibull random variable X are computed by the calculator.

Input	v (shape)	λ (scale)	γ (location)
from fit	6	100	1000

Calculated	Mean μ	Std. Dev. σ	Skew	Kurtosis
from fit	1092.7719	17.976749	-0.373262	3.0354553



ENTER the lower and upper test points (reliabilities $0 \leq p_l < p_u \leq 0.999$), fixed β risk, and the lower or upper limit on the Weibull random variable X based on the requirement.

Input	p_l	β	p_u	X_{min} or X_{max}
requirement	0.9973	0.1	0.999	1136.7

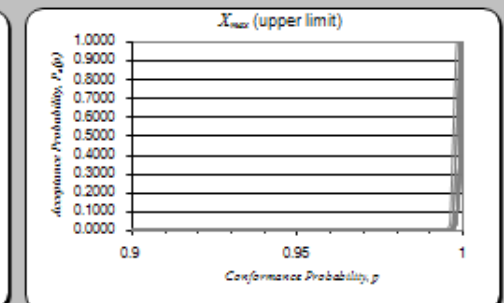
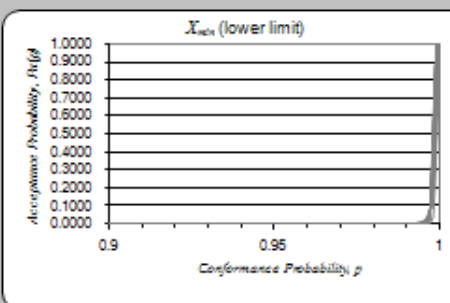
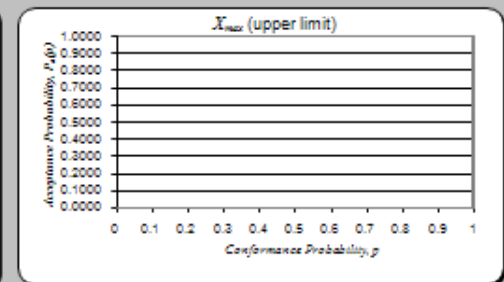
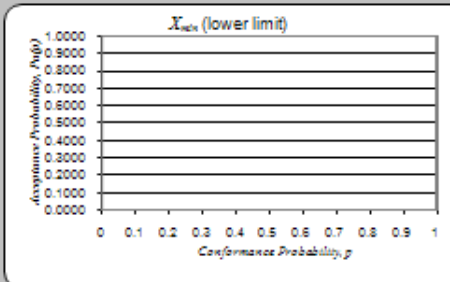
If the limit is less than the mean, the value entered is assumed to be a lower limit; if the limit is greater than the mean, the value entered is assumed to be an upper limit.

The corresponding (n, k) sampling plans and associated α risks are tabulated.

Results	n	k	α	Accept	$\mu_{weibull}$
	1486	2.3775087	0.001	Yes	1135.5118
	1348	2.380439	0.002	Yes	1135.5645
	1267	2.3823936	0.003	Yes	1135.5996
	1209	2.383911	0.004	Yes	1135.6269
	1163	2.3851735	0.005	Yes	1135.6496
	1126	2.3862671	0.006	Yes	1135.6693
	1095	2.3872393	0.007	Yes	1135.6867
	1067	2.3881198	0.008	Yes	1135.7026
	1043	2.3889281	0.009	Yes	1135.7171
	1021	2.389678	0.01	Yes	1135.7306
	876	2.3953703	0.02	Yes	1135.8329
	790	2.3994907	0.03	Yes	1135.907
	728	2.4029002	0.04	Yes	1135.9683
	679	2.4058971	0.05	Yes	1136.0221
	639	2.4086241	0.06	Yes	1136.0712
	605	2.4111618	0.07	Yes	1136.1168
	576	2.4135606	0.08	Yes	1136.1599
	549	2.4158546	0.09	Yes	1136.2011
	526	2.4180679	0.1	Yes	1136.2409
	433	2.4284349	0.15	Yes	1136.4273
	366	2.4383814	0.2	Yes	1136.6061
	313	2.4484937	0.25	No	1136.7879
	269	2.4591699	0.3	No	1136.9798
	231	2.4707793	0.35	No	1137.1885
	197	2.4837409	0.4	No	1137.4215
	168	2.4985953	0.45	No	1137.6886
	141	2.5161042	0.5	No	1138.0033

DETERMINE if the design is acceptable from the plan with the largest value of n which is no greater than the number of data points used to fit the distribution.

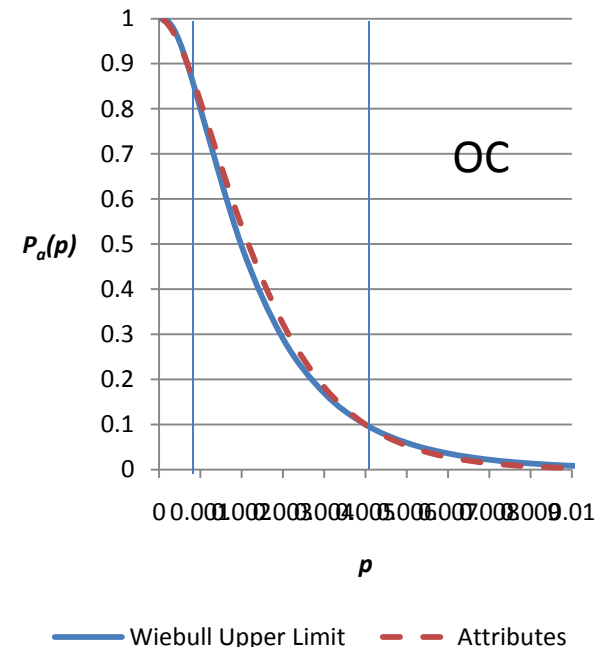
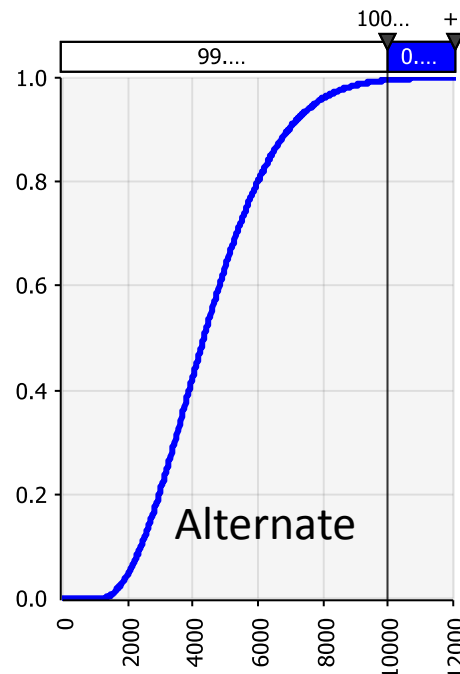
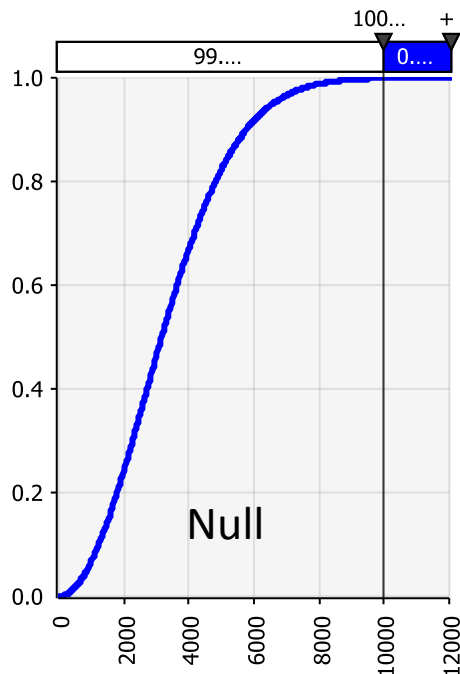
Operating Characteristic (OC) or Power Curves for all plans.



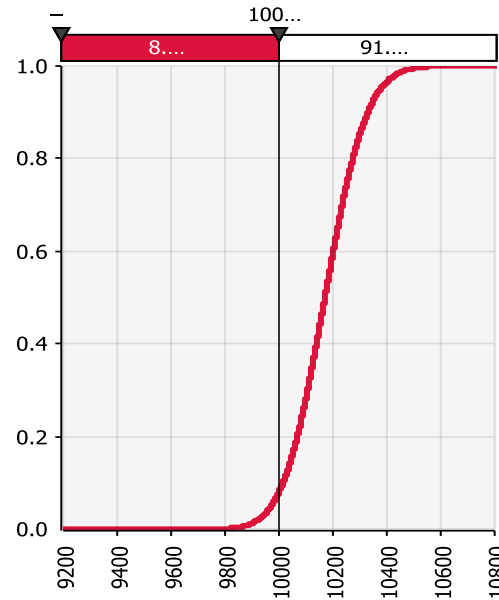
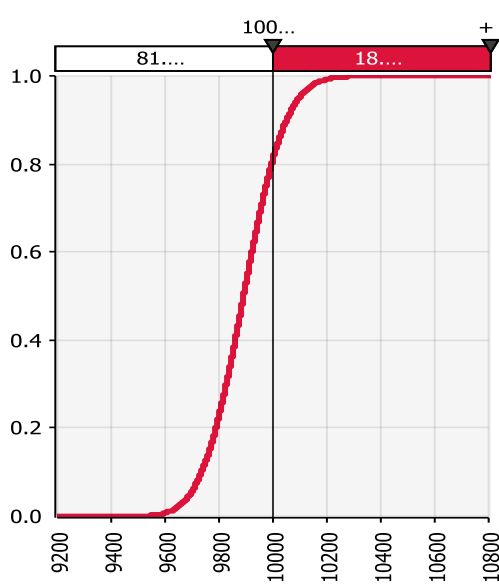
Calculations are approximate, based on the procedure given by K. Takagi (1972) "On designing unknown-sigma sampling plans based on a wide class of non-normal distributions, *Technometrics*, 14(3):669-678.

Empirical test results (example)

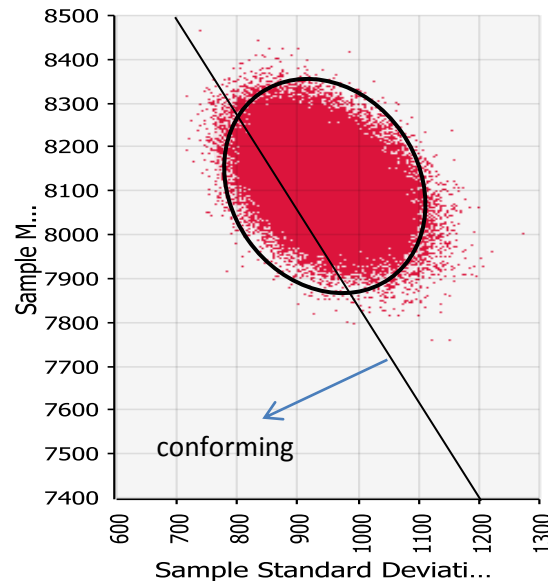
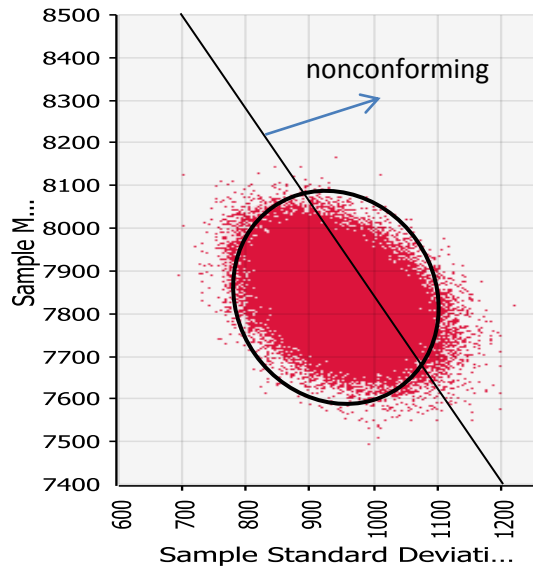
Consider an upper limit of $x_{max}=10,000$ for a random variable X distributed Weibull with unknown shift parameter δ and estimated shape and scale parameters. For the test $OC(p_0, \alpha)=(0.005, 0.2)$, $(p_1, \beta)=(0.001, 0.1)$, the associated the null and alternative means are $\mu_0=7841.64$ and $\mu_1=8121.07$, respectively. The variables plan from the gamma calculator is $(n, k)=(156, 2.17779)$.



Empirical test example results



Sampling distribution of the mean estimated using 100,000 Monte Carlo trials.



Scatter diagram for estimated μ and σ . Line is the acceptance limit

$$A = \mu - k\sigma$$

Empirical test summary results

α and β estimated using 100K Monte Carlo trials for plans with $(p_0, \alpha)=(0.005, 0.2)$, $(p_1, \beta)=(0.001, 0.1)$

Results for $x_{min}=1000$

Variable	n	k	α	β	n_v/n_a
Exponential(μ)	2	2.43×10^{-3}	0.200	0.082	0.003
Normal($\mu, \sigma=100$)	18	2.886	0.191	0.097	0.023
Normal(μ, σ)	88	2.886	0.191	0.097	0.099
Gamma(10,338 , δ)	206	2.131	0.193	0.096	0.224
Weibull(10,1995 , δ)	91	3.623	0.189	0.079	0.117
IG(1502, 100000, δ)	18	2.886	0.173	0.382	unusable

Empirical test results

Results for $x_{max}=10,000$

Variable	n	k	α	β	n_v/n_a
Exponential(μ)	66	6.26922	0.200	0.082	0.085
Gamma(10,441 , δ)	77	3.667	0.189	0.104	0.099
Weibull(10,3800 , δ)	156	3.623	0.188	0.081	0.201

Results for discrete

Variable	n	c	α	β	n_v/n_a
Binomial(n,p)	777	1	0.188	0.100	1
Poisson(n,p)	21	88	0.191	0.097	0.035

ASV fundamental assumption

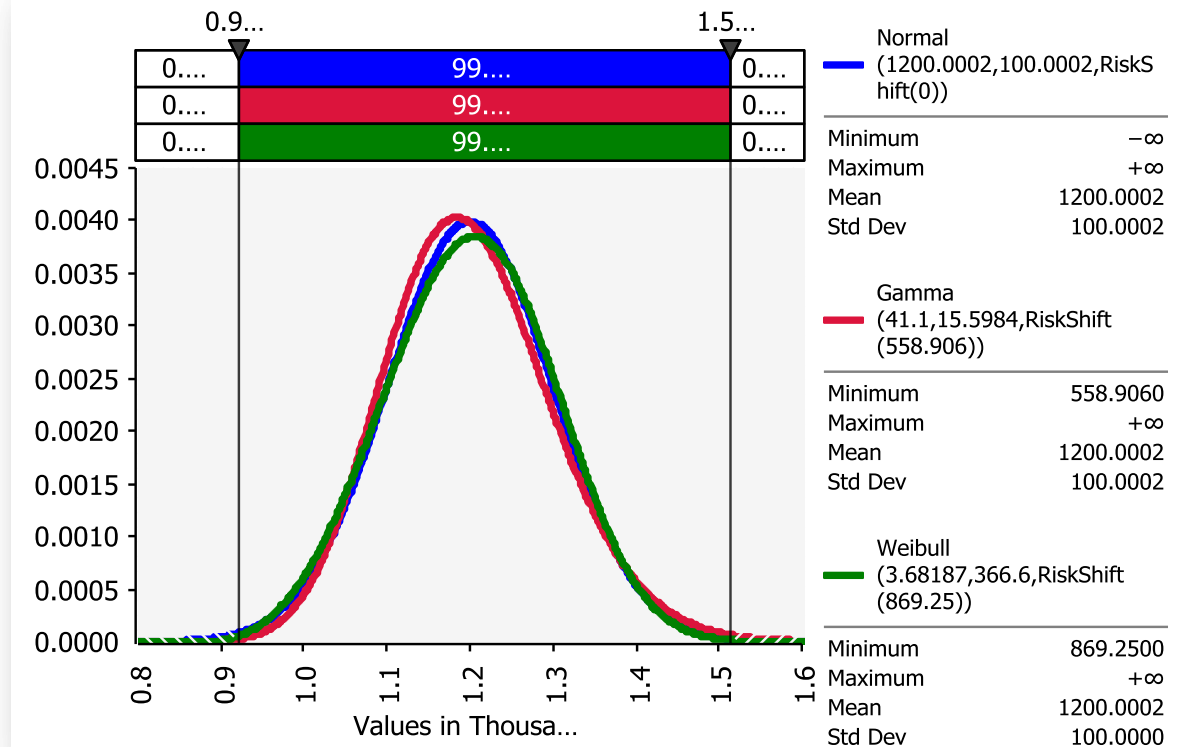
The fundamental assumption of ASV is that the **form of the output distribution is known** a priori. (Moreover, validity testing showed that ASV procedures are robust to error in the shift and scale parameters, but not to shape parameters.)

- In general, the assumption is **unsubstantiated** and the form of the distribution must be determined by **fitting** sample data.
- The question naturally arises, “How many trials are required in order to fit the correct form of the parent distribution?”
- Specifically, “**Can we obtain a unique fit to the correct parent distribution based on a sample which is approximately the same size as that specified in the corresponding ASV plan?**”
- The literature appears to be essentially silent on this issue.

For an exception, see C. Liu (1997) *A Comparison Between the Weibull and Lognormal Models Used to Analyse Reliability Data*, Ph.D. Dissertation, University of Nottingham, UK

Test case 1—Near Normal

Three parent output distributions where chosen with **similar shapes** and **identical means and standard deviations**.



Fitting tests were performed for a requirement with
condition: limit (upper or lower)

reliability: $\rho=0.9973$

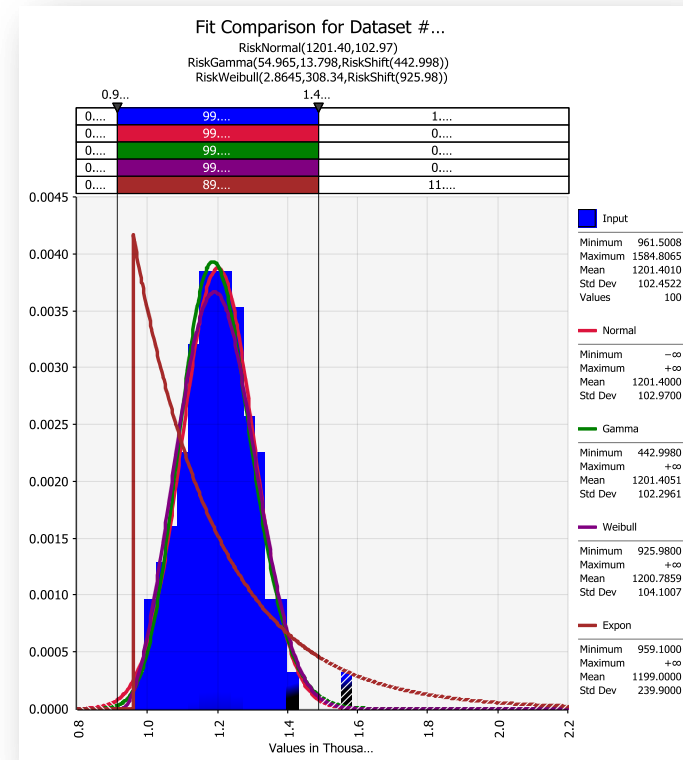
consumer's risk: $\beta=0.1$

(This seems typical of what we had been seeing as Cx level 2 requirements.)

Test 1 procedure and results

- 30 samples of 100 trials each were drawn from the parent Normal distribution.
- Normal, Gamma, Weibull, and Exponential distributions were fit to sample using commercial software (@RISK).
- In general, good fits to the Normal data were achieved with a Normal distribution.
- But good fits also were achieved with Weibull and often Gamma (but not Exponential).
- Fits were compared using three alternative goodness-of-fit (GOF) tests--the “best fit” was sensitive to the GOF test used. (Note: Anderson-Darling is the preferred test here because it gives more weight to the tails.)

- The test procedure was repeated for samples of 300 trials each, with no appreciable change in the nature of the results.
- The test procedure was repeated for samples of 100 and 300 drawn from the Gamma and Weibull parent distributions, with no appreciable change in the nature of the results.



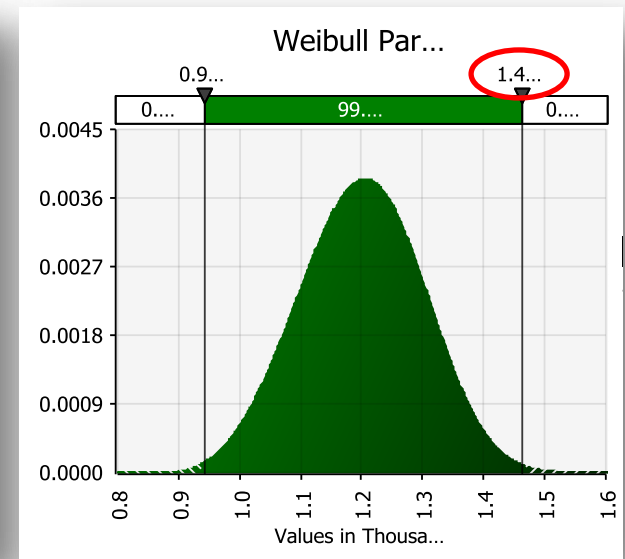
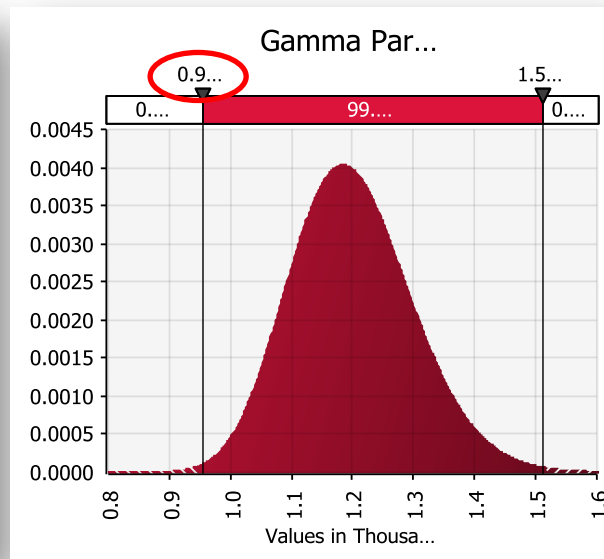
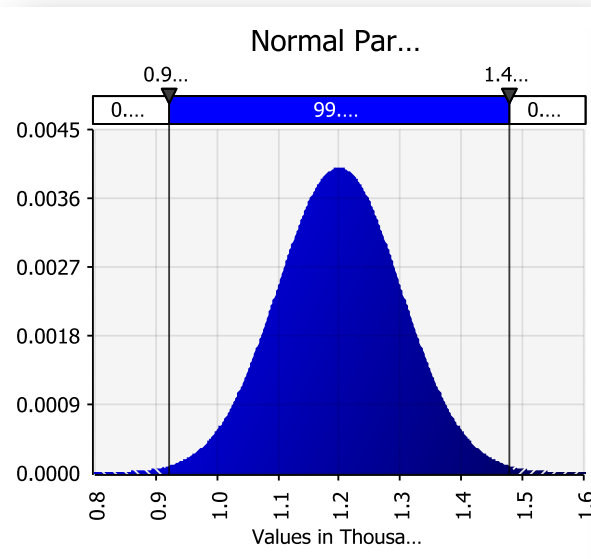
Conclusion: For our example, the size of the sampling plan is inadequate to distinguish the parent distribution for data sets with near-normal shape.

Useful result

Acceptance limits for $\rho=0.9973$ for the three parent distributions.

Parent	Lower	Upper
Normal	922	1478
Gamma	957	1513
Weibull	943	1463

- The test results were not unexpected—these distributions are very similar in shape *overall*.
- Our interest is in the (small) differences in the tails of these of these distributions.
- Note that the distribution with the **smallest k factor provides the greatest protection against accepting a nonconforming design** (i.e., the largest lower limit and the smallest upper limit).



Application

Upper limit sampling plans ($\alpha=0.2$)

Variable	n	c	k	A
Binomial (ASA plan)	2959	4		
Normal(1200, 0)	257		2.968	1497
Gamma(41.1, 15.5984, 558.906)	224		3.378	1538
Weibull(3.68187, 366.6, 869.25)	296		2.787	1479

Lower limit sampling plans ($\alpha=0.2$)

Variable	n	c	k	A
Binomial (ASA plan)	2959	4		
Normal(1200, 0)	257		2.968	903
Gamma(41.1, 15.5984, 558.906)	353		2.566	943
Weibull(3.68187, 366.6, 869.25)	615		2.677	932

The plan with the **tightest bound will yield the most conservative decision**—one that guarantees the consumer's risk is no greater than specified.

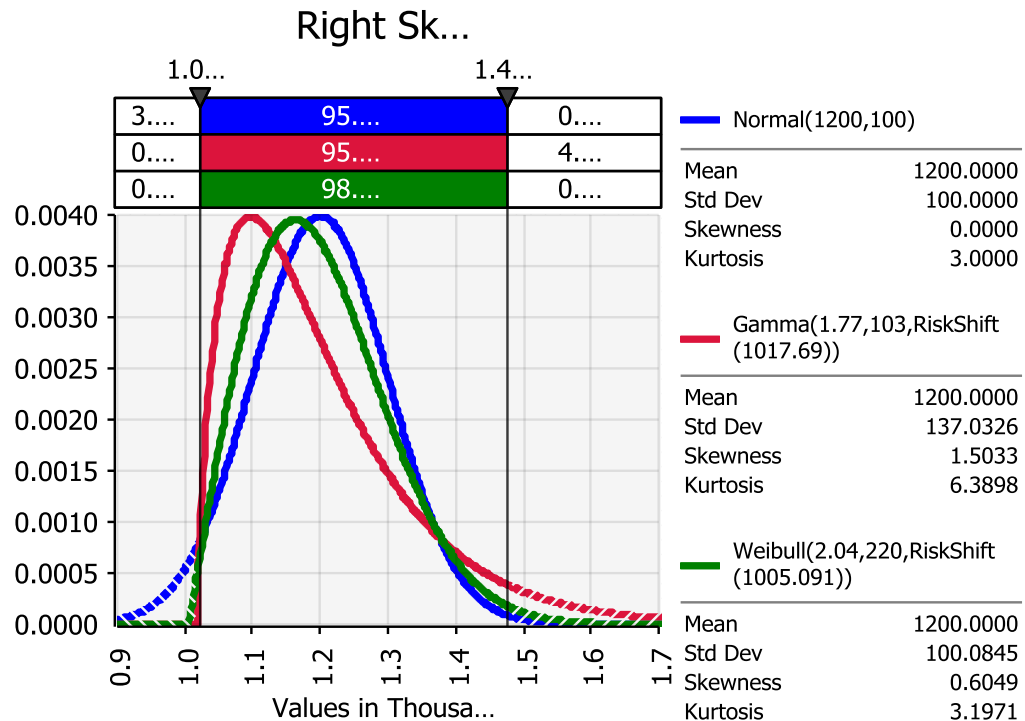
For the example requirement, the Weibull plan provides the tightest (least) upper bound.

The Gamma plan provides the tightest (greatest) lower bound.

These plans provide an order-of-magnitude reduction in computing effort.

Test case 2—Right skew

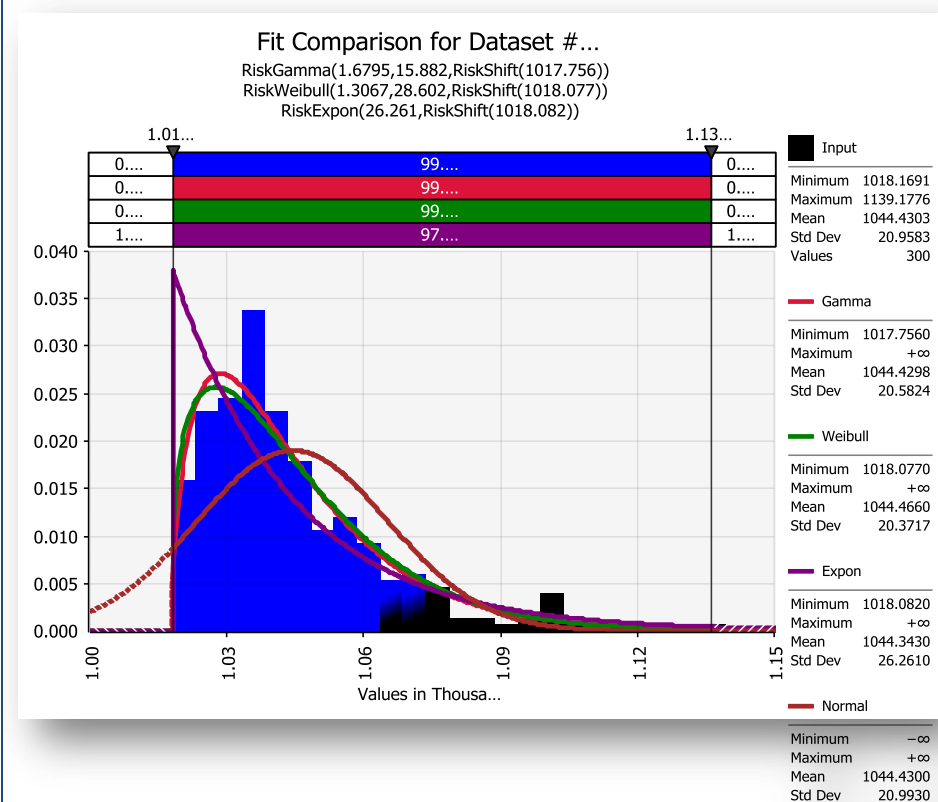
The second case considers Weibull and Gamma distributions with identical means and moderate skew.



Fitting tests were performed for the same requirement condition: limit (upper or lower)
 reliability: $\rho=0.9973$
 consumer's risk: $\beta=0.1$

Test 2 procedure and results

- 30 samples of 300 trials each were drawn from the parent Weibull distribution.
- Normal, Gamma, Weibull, and Exponential distributions were fit to sample using commercial software.
- In general, good fits to the Weibull data were achieved with the Weibull distribution and sometimes the Gamma distribution.
- The skew is sufficiently large that the data are not mistaken as Normal.
- The skew is sufficiently small that the data are not mistaken as Exponential.
- The test procedure was repeated for a the parent Gamma distribution, with no appreciable change in the nature of the results.



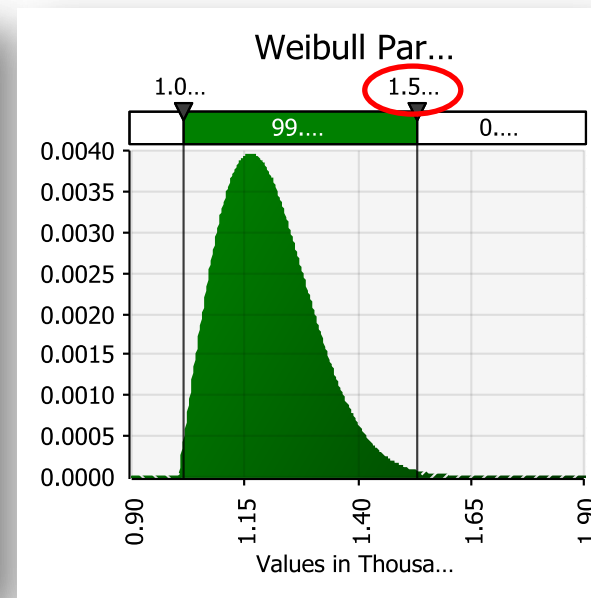
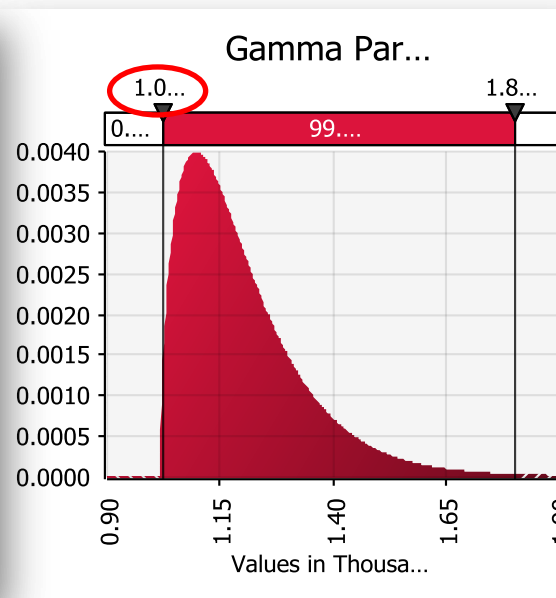
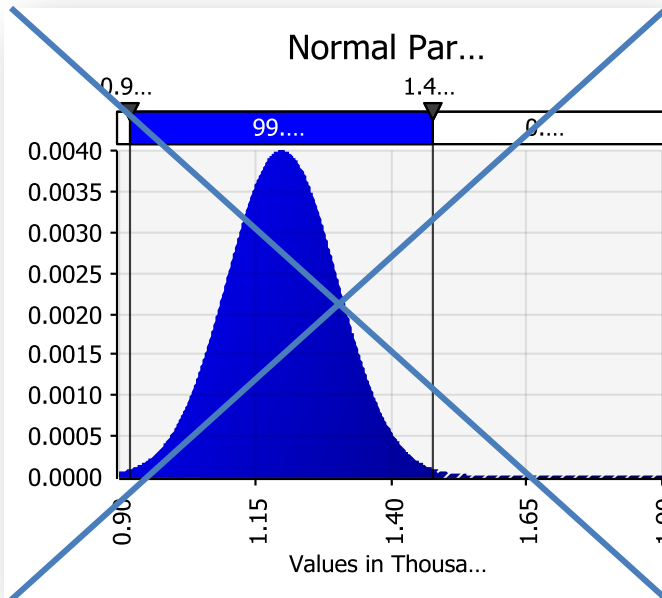
Conclusion: For our example, the size of the sampling plan is inadequate to distinguish between Weibull and Gamma, but adequate to rule out Normal and Exponential parents.

Useful result 2

Acceptance limits for $\rho=0.9973$ for the three parent distributions.

Parent	Lower	Upper
Normal	922	1478
Gamma	1023	1806
Weibull	1017	1531

- As before, the test results were not unexpected and our interest is in the (small) differences in the tails of these of these distributions.
- As before, the distribution with the **fattest tail in the direction of the limit provides the greatest protection** against accepting a nonconforming design.
- However, with sufficient skew as in this example, we can rule out Normal as the parent (poor fit).



Application

Upper limit sampling plans ($\alpha=0.2$)

Variable	n	c	k	A
Binomial (ASA plan)	2959	4		
Normal(1200, 0)	257		2.968	1497
Gamma(1.77,15.5984,1017.7)	278		4.915	1873
Weibull(2.04,220,1005.9)	264		3.557	1557

Lower limit sampling plans ($\alpha=0.2$)

Variable	n	c	k	A
Binomial (ASA plan)	2959	4		
Normal(1200, 0)	257		2.968	903
Gamma(1.77,15.5984,1017.7)	25009		1.304	1018
Weibull(2.04,220,1005.6)	3651		1.855	1015

Once again, the plan with the tightest bound will yield a conservative decision.

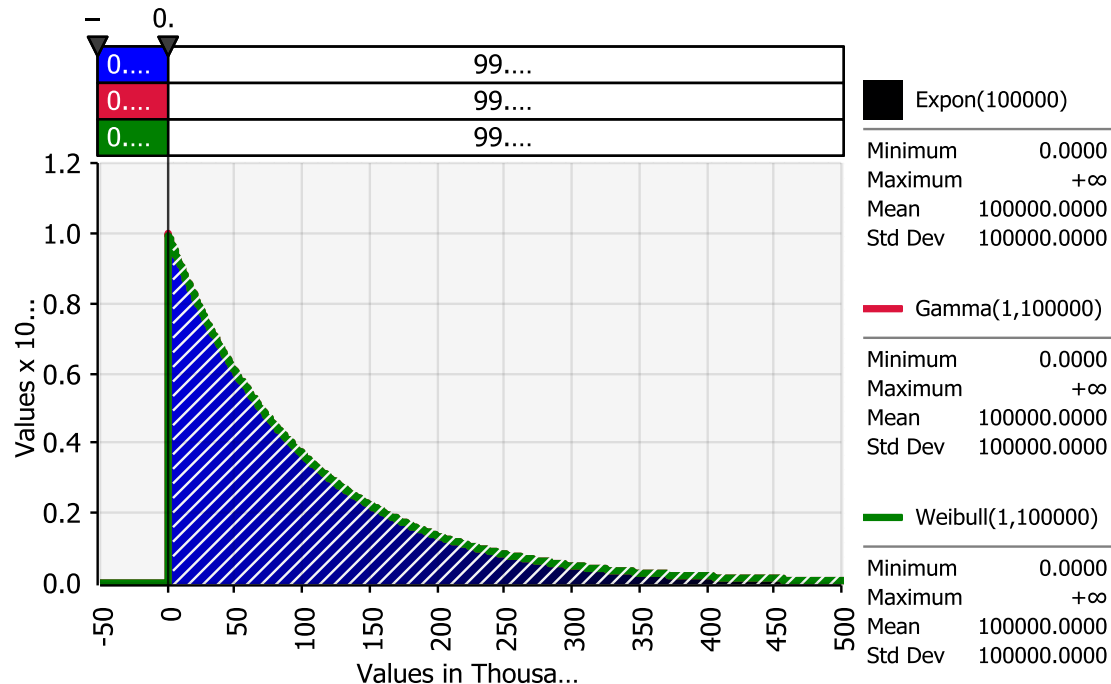
The Weibull plan provides the tightest (least) upper bound and can be used in this case.

The Gamma plan provides the tightest (greatest) upper bound, but is very large. Obviously, the ASA plan is preferable in this case.

The UL plan provides an order-of-magnitude reduction in computing effort.

Test case 3—Near Exponential LL

- A typical application is lifetime data, where the random variable X represents the time at which a component fails and the condition is the lower limit $L=X_{\min}$.
- Note that $\text{Expo}(\theta)$, $\text{Gamma}(1, \theta)$, and $\text{Weibull}(1, \theta)$ are the *same* distribution.



Plan comparisons

- The Exponential plan is *very* small—more than three orders of magnitude smaller than the attributes plan.
- The Gamma and Weibull plans are identical and *very* large—almost three orders of magnitude larger than the attributes plan.

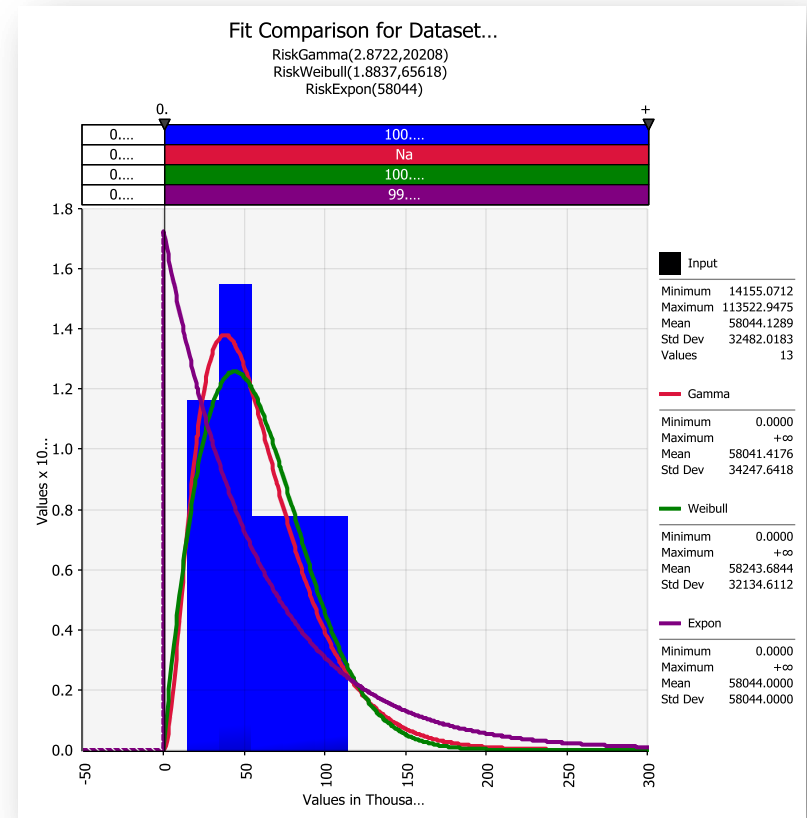
This remarkable difference in n flows from the fact that Exponential has a single parameter and fixed shape—if we *know* a priori the parent is exponential, then we need only estimate the mean.
- The Gamma and Weibull plans are the same, both derived from the same approximation. These are more conservative than the exponential (the true value of $A=270.37$), but unusable in this application because of their size.

Lower limit sampling plans ($\alpha=0.02$)

Variable	n	c	k	A
Binomial (ASA plan)	6580	12		
Exponential(100000)	13		0.9806	1940.00
Gamma(1, 100000)	3819299		0.9979	2060.45
Weibull(1,100000)	3819299		0.9979	2060.45

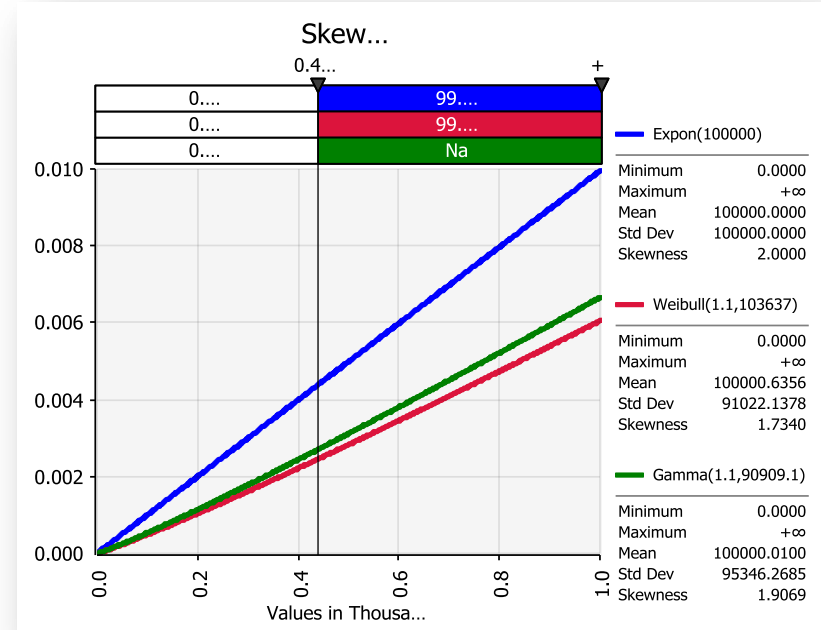
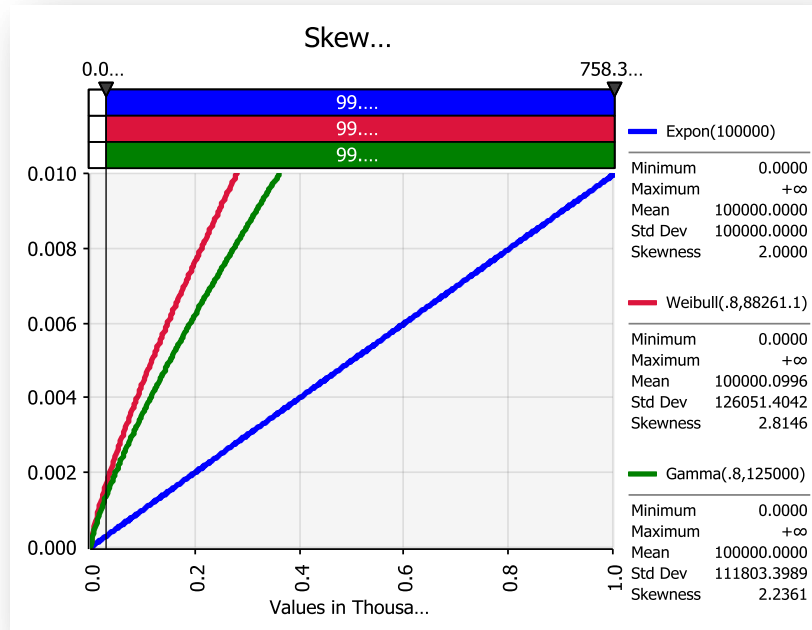
Test case 3—modest skew

- 30 samples of 13 trials each were drawn from the parent Exponential distribution
- Gamma, Weibull, and Exponential distributions were fit to these datasets (shifts set to zero for lifetime data).
- Fits were compared using three alternative goodness-of-fit tests.
- Best fits were dependent GOF test and there were many ties for best fit.
- Acceptable fits to exponential ($p\text{-value} \geq 0.15$) were obtained in most cases (30 Chi-squared, 25 K-S, 23 A-D).
- Exponential or binomial are the only practical plans in this case and acceptable Exponential fits most often can be achieved for Exponential parent distributions.



Effect of skew

But what about datasets from Weibull and Gamma parents masquerading as Exponential?



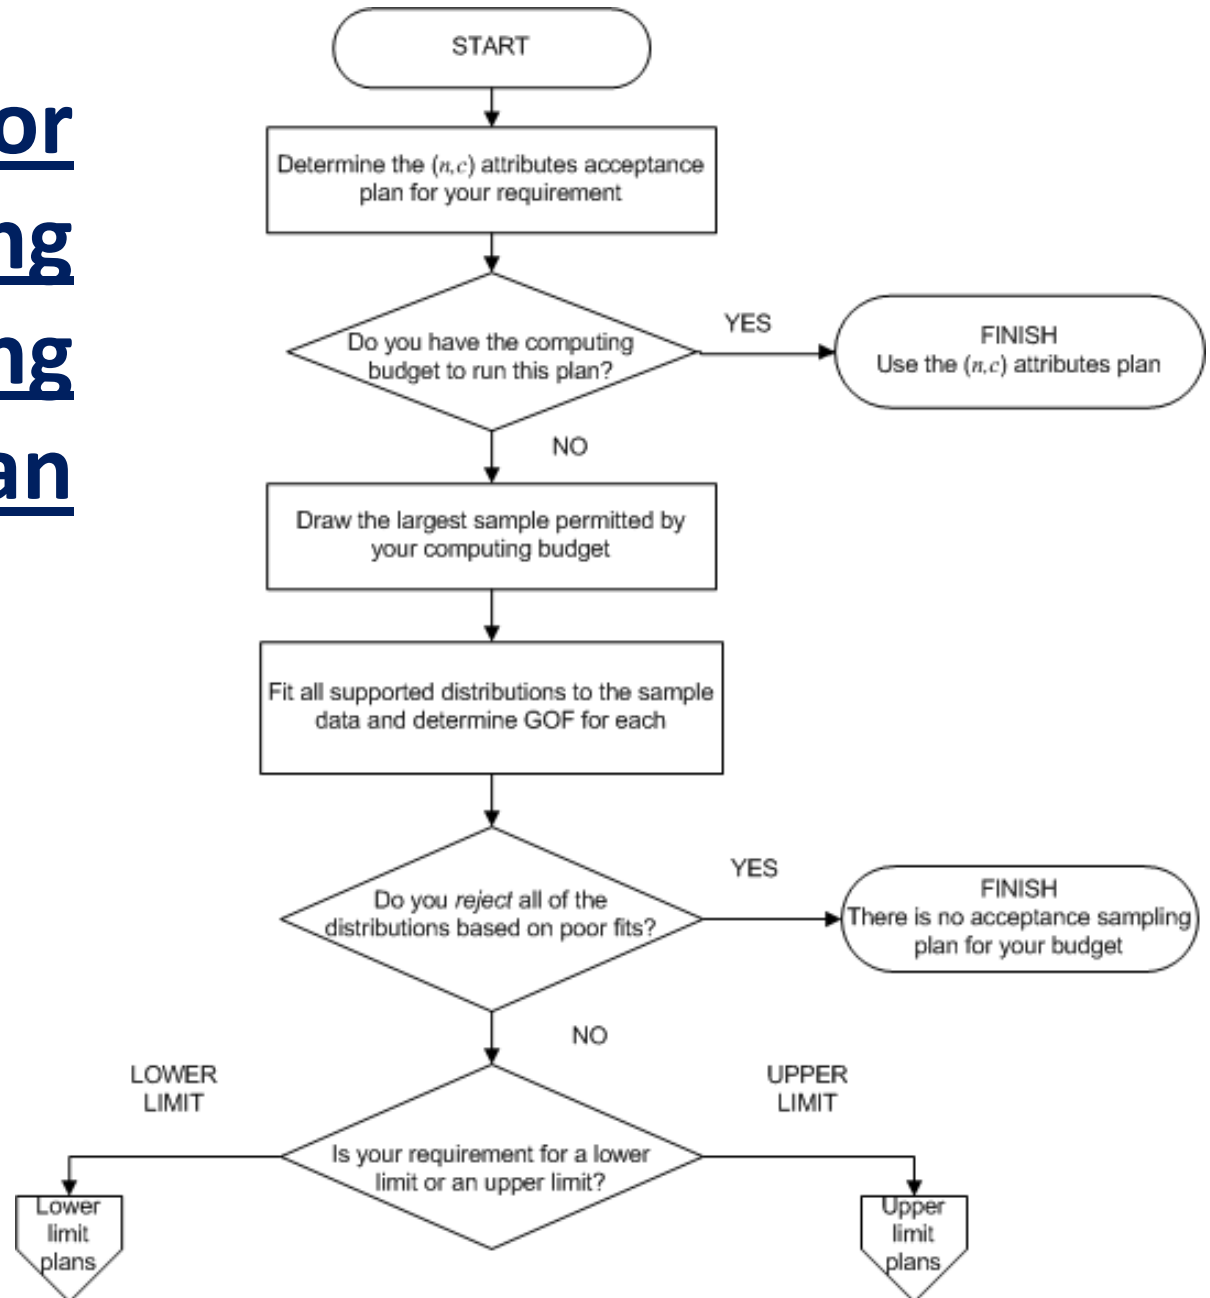
The Exponential plans are conservative for Weibull parents with skew >2 and non-conservative for skew <2 . Accepting the Exponential fit in the second case will result in a modestly lower reliability than specified (0.995 rather than 0.997 in this example).

Can modest skew be detected?

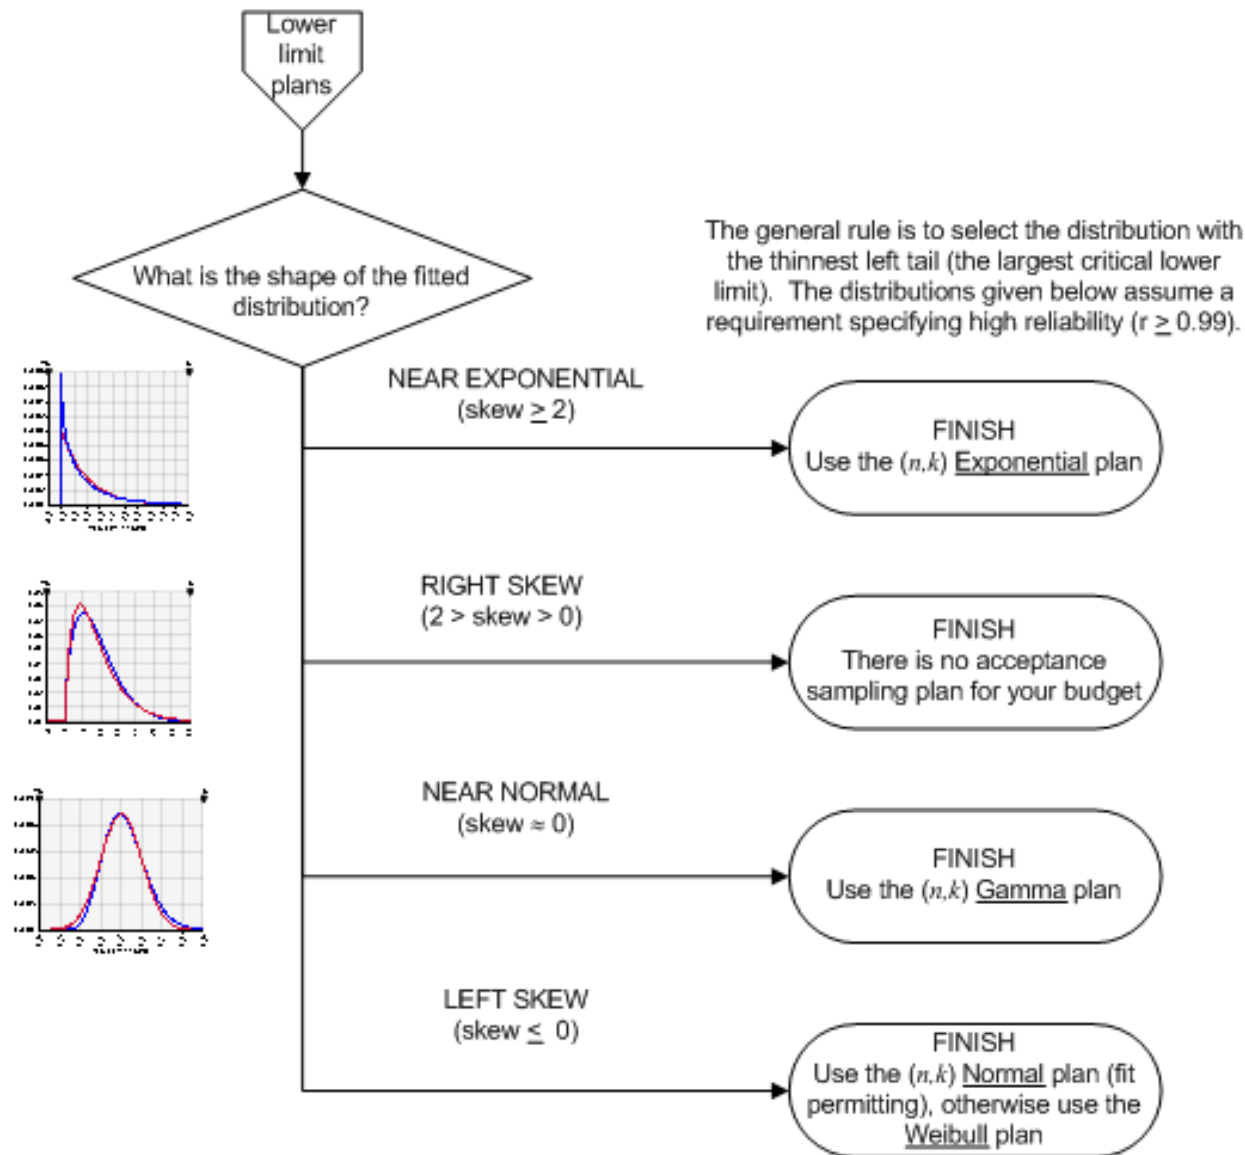
- 30 samples of 13 trials each were drawn from the parent Weibull distribution with skew=0.631
- Gamma, Weibull, and Exponential distributions were fit to these datasets (shifts set to zero for lifetime data).
- A-D and K-S tests typically showed very poor fits to Exponential even with these small samples (Chi-squared test appears to lack power).
- P-P and Q-Q plots illustrated that the Exponential fits were poorest in the region of interest .

Conclusion: With a small sample and skew ≈ 2 it is not possible to discern the parent distribution; with lesser skew Weibull and Gamma will not be confused with Exponential.

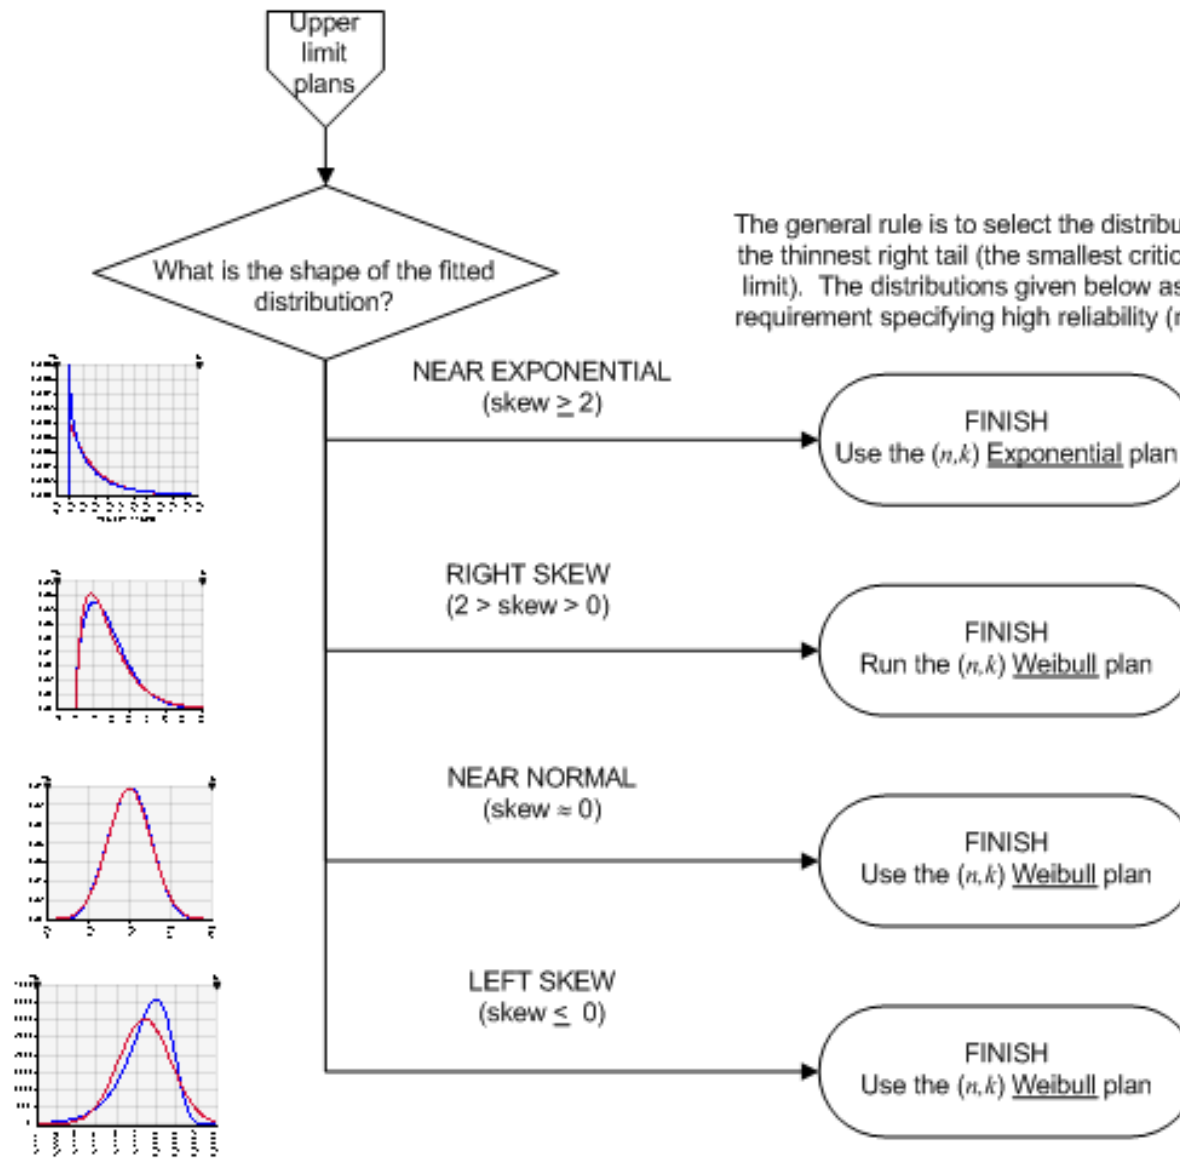
Procedure for determining a sampling plan



Procedure (lower limit plans)



Procedure (upper limit plans)



Summary

- ASA plans are preferred when computational demands can be met. These plans are exact and transparent.
- ASV plans are a viable alternative, when ASA plans are too large. These plans are inherently approximate; the data (perhaps transformed) must fit the a distribution for which ASA calculator is available. The assistance of a statistician would be beneficial.
- For data with skew less than Exponential, the Normal, Gamma, or Weibull plan with the tightest bound is a good choice—it is conservative and can provide an order of magnitude reduction in computational effort.
- For near-Exponential data with a lower limit, the Exponential plan is not necessarily conservative (and should be applied intelligently and with caution)—but can provide several orders of magnitude reduction in computational effort.
- In some applications, uncertainties in the protection afforded by Exponential may be inconsequential when model error is considered.
- When applying the Exponential, a good practice is to make as many trials as feasible and then attempt a fit to the distributions currently supported.

Contributions

Variables acceptance can reduce sample sizes and the reduction can be as much as one, two, and even three orders of magnitude depending on the distribution and OC. But this isn't always the case. Gamma and Weibull plans become larger than attributes plans as the shape parameter decreases.

Normal plans don't work well for inverse Gaussian. But the published inverse Gaussian plans don't work either. We've found the error in the derivation and I think I may have a fix given the time to mess with it. That's news.

The purpose of variables acceptance is to reduce sample size. But in accomplishing this we can't be sure that we satisfy the fundamental assumption that the distribution is known, at least for the OC we are interested in. Assuming we want to be conservative with respect to consumer's risk, we've developed a procedure to overcome this issue.